arXiv:1412.2497v1 [cs.CC] 8 Dec 2014

Homonym Population Protocols, or Providing a Small Space of Computation Using a Few Identifiers

Olivier BOURNEZ[1], Johanne COHEN[2], and Mikaël RABIE[1]

1 **LIX, Ecole Polytechnique, 91128 Palaiseau Cedex, France**
`first.last-name@lix.polytechnique.fr`
2 **LRI, Université Paris Sud, 91460 Cedex, France** `Johanne.Cohen@lri.fr`

Abstract

Population protocols have been introduced by Angluin *et al.* as a model in which n passively mobile anonymous finite-state agents stably compute a predicate on the multiset of their inputs via interactions by pairs. The model has been extended by Guerraoui and Ruppert to yield the community protocol models where agents have unique identifiers but may only store a finite number of the identifiers they already heard about. The population protocol models can only compute semi-linear predicates, whereas in the community protocol model the whole community of agents provides collectively the power of a Turing machine with a $O(n \log n)$ space.

We consider variations on the above models and we obtain a whole landscape that covers and extends already known results: By considering the case of homonyms, that is to say the case when several agents may share the same identifier, we provide a hierarchy that goes from the case of no identifier (i.e. a single one for all, i.e. the population protocol model) to the case of unique identifiers (i.e. community protocol model).

We obtain in particular that any Turing Machine on space $O(\log^{O(1)} n)$ can be simulated with at least $O(\log^{O(1)} n)$ identifiers, a result filling a gap left open in all previous studies.

Our results also extend and revisit in particular the hierarchy provided by Chatzigiannakis *et al.* on population protocols carrying Turing Machines on limited space, solving the problem of the gap left by this work between per-agent space $o(\log \log n)$ (proved to be equivalent to population protocols) and $O(\log n)$ (proved to be equivalent to Turing machines).

1 Introduction

Angluin *et al.* [**?**] proposed a model of distributed computation called *population protocols*. Population protocols can be seen as a minimal model that aims at modeling large sensor networks with resource-limited anonymous mobile agents. The mobility of the agents is assumed to be unpredictable (given by any fair scheduler) and pairs of agents can exchange state information when they are close together.

The population protocol model can be considered as a computational model, in particular computing predicates: Given some input configuration, the agents have to decide whether this input satisfies the predicate. More precisely, the population of agents has to eventually stabilize to a configuration in which every agent is in an accepting state or a rejecting one. This must happen with the same program for all population sizes, i.e for any size of input configuration.

The seminal work of Angluin *et al.* [**?**, **?**] proved that predicates computed by population protocols are precisely those on counts of agents definable by a first-order formula in Presburger arithmetic (or equivalently must correspond to a semilinear set). Subsets definable in this way are rather restricted, as for example multiplication is not expressible in Presburger arithmetic. Several variants of the original model have been considered in order to

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

strengthen the population protocol model with extra realistic and implementable assumptions, in order to gain more computational power. Variants also include natural restrictions like modifying the assumption between agent's interactions (one-way communications [**?**], particular interaction graphs [**?**]). This also includes the Probabilistic Population Protocol model that makes a random scheduling assumption for interactions [**?**]. Various kinds of fault tolerance have been considered for population protocols [**?**], including the search for self-stabilizing solutions [**?**]. We refer to [**?**, **?**, **?**] for a survey.

Among many variants of population protocols, the *passively mobile (logarithmic space) machine model* introduced by Chatzigiannakis *et al.* [**?**] constitutes a generalization of the population protocol model where finite state agents are replaced by agents that correspond to arbitrary Turing machines with $O(S(n))$ space per-agent, where n is the number of agents. An exact characterization [**?**] of computable predicates is given: this model can compute all symmetric predicates in $NSPACE(nS(n))$ as long as $S(n) = \Omega(\log n)$. Chatzigiannakis *et al.* establish that with a space in agent in $S(n) = o(\log \log n)$, the model is equivalent to population protocols, i.e. to the case $S(n) = O(1)$.

In parallel, *community protocols* introduced by Guerraoui and Ruppert [**?**] are closer to the original population protocol model, assuming *a priori* agents with individual very restricted computational capabilities. In this model, each agent has a unique identifier. Each agent can only remember $O(1)$ other agent identifiers, and only identifiers of agents that it met. Guerraoui and Ruppert [**?**], using results about the so-called storage modification machines [**?**], proved that such protocols can actually still simulate collectively a Turing machine: Predicates computed by this model with n agents are precisely the predicates in $NSPACE(n \log n)$.

In this paper, we obtain a whole landscape that covers and extends already known results. We do so by considering that the capabilities of agents is even more restricted.

Indeed, we drop the hypothesis of unique identifiers. That is to say, we consider that agents may have homonyms. We obtain a hierarchy that goes from the case of no identifier (i.e. the same identifier for everyone, i.e. the population protocol model) to the case of unique identifiers (i.e. community protocol model). In what follows, $f(n)$ denotes the number of available identifiers on a population with n agents. Notice that the idea of having less identifiers than agents, that is to say of having "homonyms", has already been considered in other contexts or with not closely related problematics [**?**, **?**, **?**].

Basically, our results are summarized in Tables 1 and 2 where $MNSPACE(S(n))$ (respectively: $SMNSPACE(S(n))$) is the set of f-symmetric[1] (resp. also stable under the permutation of the input multisets) languages recognized by Non Deterministic Turing Machines on space $O(S(n))$.

Our results can then also be considered as extending the *passively mobile machine model* of Chatzigiannakis *et al.* [**?**]. We solve in particular the problem of the gap left between space per agent $S(n) = o(\log \log n)$ (proved to be equivalent to population protocols) and $S(n) = O(\log n)$ (proved to be equivalent to Turing machines): With $S(n) = O(\log \log n)$, the model is equivalent to $\bigcup_k NSPACE(\log^k n)$ (see Table 2).

The paper is organized as follows: Section 2 introduces the formal definitions of the different models and some already known main results. Section 3 is devoted to the case when we assume that an order is available on identifiers, but that the number $f(n)$ of

[1] These elements are defined in Section 3.4.

$f(n)$ identifiers	Computational power
$O(1)$	Presburger's definable subsets [?, ?]
$\Theta(\log^r n)$ $r \in \mathbb{R}_{>0}$	$\bigcup_{k\in\mathbb{N}} MNSPACE\left(\log^k n\right)$ Theorem 10
$\Theta(\sqrt[k]{n})$ with $k \in \mathbb{N}_{>0}$	$MNSPACE(n \log n)$ Theorem 11
n	$NSPACE(n \log n)$ [?]

Table 1 Homonym population protocols with n agents and $f(n)$ identifiers.

Space per agent $S(n)$	Computational power
$O(1)$	Presburger's definable subsets [?, ?]
$o(\log \log n)$	Presburger's definable subsets [?]
$\Theta(\log \log n)$	$\bigcup_{k\in\mathbb{N}} SNSPACE(\log^k n)$ Theorem 13
$\Omega(\log n)$	$SNSPACE(nS(n))$ [?]

Table 2 Passively mobile machine model [?] with n agents and space $S(n)$ by agent.

identifiers is possibly less than the number n of agents (see Table 1). Section 4 treats the case $O(\log \log n)$ of the passively mobile machines [?] (see Table 2). Section 5 is then a summary of our results with some open questions.

2 Models

The models we consider are basically variations of the community protocol model from [?]. This latter model can in turn be considered as an extension of the model from Angluin et al. [?, ?]. In all these models, a system is a collection of agents. Each agent has a finite number of possible states. Each agent has an input value, that determines its initial state. Evolution of states of agents is the product of pairwise interactions between agents: when two agents meet, they exchange information about their states and simultaneously update their own states according to a joint transition function, which corresponds to the algorithm of the protocol. The precise sequence of agents involved under the pairwise interactions is under the control of any fair scheduler. In other words, agents are passively mobile, using the terminology from [?, ?]. The considered notion of fairness for population protocols states that every system configuration that can be reached infinitely often is eventually reached.

Population protocols have been, to date, mostly considered as computing predicates: one considers protocols such that starting from some initial configuration, any fair sequence of pairwise interactions must eventually yield to a state where all agents agree and either accept or reject. The corresponding predicate then corresponds to the inputs that eventually lead to accept. Algorithms are assumed to be uniform: the protocol descriptions are independent of the number n of the agents.

The original model [**?**, **?**] has been extended in various ways: in particular, it has been extended by Chatzigiannakis *et al.* [**?**] to yield the *passively mobile machine model*. In this latter model, agents are no longer finite-state machines but are arbitrary $S(n)$-space Turing machines. We will write $PMSPACE(S(n))$ for the problems solvable in this model.

We will mainly focus on another extension of the original model proposed by Guerraoui and Ruppert [**?**]: The *community protocol* model assigns a unique *identifier* to agents, and not arbitrary Turing machines. To avoid multiplication of names, we will write community protocol for the model of [**?**], and homonym population protocol with $f(n)$ distinct identifiers for our version.

Let U be the infinite set containing the possible *id*s (identifiers). Compared to [**?**], we do not consider that the set is arbitrary: we assume that $U \subset \mathbb{N}$. We also assume these *identifier*s not to be unique: several agents may have the same *identifier*. We only assume that there are $f(n)$ possibles identifiers.

More formally, a community protocol / homonym population protocol algorithm is then specified by:

1. a function f associating to the size of the population the number of *identifiers* appearing in this population;
2. a finite set B of possible basic states;
3. an integer $d \geq 0$ representing the number of *id*s that can be remembered by an agent;
4. some input alphabet Σ and some output alphabet Y;
5. an input map $\iota : \Sigma \to B$ and an output map $\omega : B \to Y$;
6. a transition function $\delta : Q^2 \to Q^2$, with $Q = B \times U \times (U \cup \{_\})^d$.

▶ Remark. We assume to simplify writing in the following that δ is a function, but this could be a relation as in [**?**], without changing our results.

The state of an agent stores an element of B, together with up to d *id*s. The state of an agent is assumed to be initialized with some *id*. If any of the d slots is not currently storing an *id*, it contains the *null id* $_ \notin U$. In other words, $Q = B \times (U \cup \{_\})^d$ is the set of possible agent states. The transition relation indicates the result of a pairwise interaction: when agents in respective state q_1 and q_2 meet, they move to respectively state q_1' and q_2' whenever $\delta(q_1, q_2) = (q_1', q_2')$.

As in [**?**], we assume that agents store only *id*s they have learned from other agents (otherwise, they could be used as an external way of storing arbitrary information and this could be used as a space for computation in a non interesting and trivial way): we assume that if $\delta(q_1, q_2) = (q_1', q_2')$, and id appears in q_1', q_2' then id must appear in q_1 or in q_2.

As in [**?**], we assume that the identifiers of agents are chosen by some adversary, and not under the control of the program.

We add an hypothesis to the model of [**?**]: agents need to know when an *id* is equal to 0 and when an *id* is the direct successor of another one. As we want to be minimal, we hence assume that this is the only hypothesis we make on *id*s in the following section. More formally, whenever $\delta(q_1, q_2) = (q_1', q_2')$, let $u_1 < u_2 < \cdots < u_d$ be the distinct *id*s that appear in any of the four states q_1, q_2, q_1', q_2'. Let $v_1 < v_2 < \cdots < v_k$ such that $u_1 = 0 \Leftrightarrow v_1 = 0$ and $v_i + 1 = v_{i+1} \Leftrightarrow u_i + 1 = u_{i+1}$. If $\rho(q)$ is the state obtained from q by replacing all occurrences of each *id* u_i by v_i, then we require that $\delta(\rho(q_1), \rho(q_2)) = (\rho(q_1'), \rho(q_2'))$.

▶ Remark. This weakening of the original model does not change the computational power in the case where all agents have distinct identifiers.

▶ Remark. Our purpose is to establish results with minimal hypotheses. Our results work basically when *id*s are consecutive integers, say $\{0, 1, 2, \ldots, f(n) - 1\}$. This may be thought

as a restriction. This is why we weaken to the above hypothesis, which seems to be the minimal hypothesis to make our proofs and constructions correct.

Without this hypothesis that we think "minimal" (i.e. without the possibility to know if an *id* is the successor of another one) we think that the model is far too weak.

▶ Remark. Notice that knowing if an *id* is equal to 0 is not essential, but ease the explanation of our counting protocols.

A configuration of the algorithm then consists of a finite vector of elements from Q. An input of size $n \geq 2$ is $f(n)$ non empty multisets X_i over alphabet Σ (corresponding to the $f(n)$ identifiers that we assume to be $0, 1, \ldots, f(n)-1$ in order to ease the explanation[2]). An initial configuration for n agents is a vector in Q^n of the form $((\iota(x_j), i-1, _, \ldots, _))_{1 \leq i \leq f(n), 1 \leq j \leq |X_i|}$ where x_j is the jth element of X_i: in other words, every agent starts in a basic state encoding $\iota(x_j)$, its associated identifier and no other identifier stored in its d slots.

If $C = (q_1, q_2, \ldots, q_n)$ and $C' = (q'_1, q'_2, \ldots, q'_n)$ are two configurations, then we say that $C \to C'$ (C' is reachable from C in a unique step) if there are indices $i \neq j$ such that $\delta(q_i, q_j) = (q'_i, q'_j)$ and $q'_k = q_k$ for all k different from i and j. An execution is a sequence of configurations $C_0, C_1, \ldots$, such that C_0 is an initial configuration, and $C_i \to C_{i+1}$ for all i. An execution is fair if for each configuration C that appears infinitely often and for each C' such that $C \to C'$, C' appears infinitely often.

▶ **Example 1** (Leader Election)**.** We adapt here a classical example of Population Protocol. We want a protocol that performs a leader election, with the additional hypothesis that when the election is terminated, all agents know the identifier of the leader (for the classical Population Protocol, it is not possible to store the identifier of the leader). If one prefers, each agent with identifier k starts with state $L_{k,_}$, considering that it is a leader, with identifier k. We want that eventually at some time (i.e in a finite number of steps), there will be a unique agent in state L_{k_0,k_0}, where k_0 is the identifier of this unique agent, and all the other agents in state N_{i,k_0} (where i is its *id*).

A protocol that solves the problem is the following (from now on, we will denote by $q_{k,k'}$ an agent in state q with *id* k storing the second *id* q', as $d = 1$): $f(n) = n$, $B = \{L, N\}$, $d = 1$ (only the identifier of the current leader is stored), $\Sigma = \{L\}$, $Y = True$, $\iota(L) = L$, $\omega(L) = \omega(N) = True$, and δ such that the rules are:

$$
\begin{array}{llll}
L_{k,_} & q & \to L_{k,k} \quad q & \forall k \in \mathbb{N}, \forall q \in Q \\
L_{k,k} & L_{k',k'} & \to L_{k,k} \quad N_{k',k} & \forall k, k' \\
L_{k,k} & N_{i,k'} & \to L_{k,k} \quad N_{i,k} & \forall k, k', i \\
N_{i,k'} & L_{k,k} & \to N_{i,k} \quad L_{k,k} & \forall k, k', i \\
N_{i,k} & N_{i',k'} & \to N_{i,k} \quad N_{i',k'} & \forall k, k', i, i'
\end{array}
$$

By the fairness assumption, this protocol will reach a configuration where there is exactly one agent in state L_{k_0,k_0} for some identifier k_0. Then, by fairness again, this protocol will reach the final configuration $L_{k_0,k_0} \bigcup_{i \neq k_0} N_{i,k_0}$.

A configuration has an *Interpretation* $y \in Y$ if, for each agent in the population, its state q is such that $\omega(q) = y$. If there are two agents in state q_1 and q_2 such that $\omega(q_1) \neq \omega(q_2)$, then we say that the configuration has *No Interpretation*. A protocol is said to *compute the output* y from an input x if, for each fair sequence $(C_i)_{i \in \mathbb{N}}$ starting from an initial condition

[2] Algorithms can easily be adapted if this is not the case.

C_0 representing x, there exists i such that, for each $j \geq i$, C_j has the interpretation y. The protocol is said to compute function h if it computes $y = h(x)$ for all inputs x. A predicate is a function h whose range is $Y = \{0, 1\}$.

Observe that population protocols of [**?**, **?**] are the special case of the protocols considered here where $d = 0$ and $f(n) = 1$. The following is known for the original model [**?**, **?**]:

▶ **Theorem 2** (Population Protocols [**?**]). *Any predicate over $\mathbb{N}^k$ that is first order definable in Presburger's arithmetic can be computed by a population protocol. Conversely, any predicate computed by a population protocol is a subset of $\mathbb{N}^k$ first order definable in Presburger's arithmetic.*

For the community protocols, Guerraoui and Ruppert established in [**?**] that computable predicates are exactly those of $NSPACE(n \log n)$, i.e. those in the class of languages recognized in non-deterministic space $n \log n$.

Notice that their convention of input in [**?**] requires that the input be distributed on agents ordered by identifiers.

▶ **Theorem 3** (Community Protocols [**?**]). *Any predicate in $NSPACE(n \log n)$ can be computed by a community protocol. Conversely, any predicate computed by such a community protocol is in the class $NSPACE(n \log n)$.*

Notice that Guerraoui and Ruppert [**?**] established that this holds even with Byzantine agents, under some rather strong conditions. We now determine what can be computed when the number of identifiers $f(n)$ is smaller than n. This will be done by first considering some basic protocols.

3 When Identifiers are Missing

3.1 Computing the size of the population

We first construct a protocol that computes n, that is to say the size of the population. Of course, since agents have a finite state, no single agent can store the whole size. We mean by "that computes n", the fact that the size of the population will be encoded by the global population.

Indeed, the protocol will perform a leader election over each identifier. We will call the set of leaders a *chain*. The size of the population will be written in binary on this chain (it will be possible as $f(n) \geq \log n$ in this part).

Clearly, once such a chain has been constructed, it can be used to store numbers or words, and can be used as the tape of a Turing Machine: We will (often implicitly) use in our description this trick in what follows. This will be used to simulate Turing machines in an even trickier way, in order to reduce space or identifiers.

▶ Proposition 1 (Counting Protocol). When we have $f(n)$ identifiers with $f(n) \geq \log n$, there exists an homonym population protocol that computes n: At some point, there are exactly $f(n)$ agents not in a particular state $\perp$, all having distinct identifiers. If we align these agents from the highest *id* to the lowest one, we get n written in binary.

▶ Remark. Understand that at that point, no agent knows the value of n (nor that the computation is over as usual for population protocol models). However, at that point, collectively the population encodes n.

Proof. Informally, the protocol initializes all agents to a particular state A. In parallel, it performs a leader election inside subsets of agents with same *id*. It also counts the number

of agents: an agent that has already been counted is marked in a state different from A, and will not be used in the protocol anymore; An agent in state 1 (respectively 0, or 2) with the *id* k represents 2^k (respectively 0, or 2^{k+1}) agents counted. Interactions between agents update those counts.

More formally, here is the protocol. We denote by q_k an agent in state q with Id k. The rules are as follows:

A_0	$q_k \to 1_0$	q_k	$\forall q, k$	0_k	1_k	$\to \perp_k$	1_k	$\forall k$
A_k	$0_0 \to 0_k$	1_0	$\forall k \geq 1$	1_k	1_k	$\to \perp_k$	2_k	$\forall k$
A_k	$1_0 \to 0_k$	2_0	$\forall k \geq 1$	2_k	0_{k+1}	$\to 0_k$	1_{k+1}	$\forall k$
0_k	$0_k \to \perp_k$	0_k	$\forall k$	2_k	1_{k+1}	$\to 0_k$	2_{k+1}	$\forall k$

This protocol has 3 steps. 1) At the beginning, all agents are in state A. A state A is transformed into a state 1, by adding 1 to an agent of identifier 0 (the 3 first rules). 2) Rules 4 to 6 perform a leader election for each identifier, by merging the counted agents. 3) The remaining rules correspond to summing together the counted agents, carrying on to the next identifier the 1.

Let v be the function over the states defined as follows for any k: $v(A_k) = 1$, $v(0_k) = v(\perp_k) = 0$, $v(1_k) = 2^k$, $v(2_k) = 2^{k+1}$. We can notice that the sum (of v values) over the agents remains constant over the rules. Thus the sum always equals the number of agents in the population. By fairness, this protocol reaches the desired end. ◀

▶ Remark. The previous counting protocol also works with $f(n) = \Omega(\log n)$: If $f(n) \geq \alpha \log n$ with $\alpha < 1$, then, using a base $\lceil e^{1/\alpha} \rceil$ instead of a base 2 ensures that n can be written on $f(n)$ digits.

▶ Remark. We use here the fact that the population knows if an *id* is equal to 0. We can work with identifiers in $[a, a+f(a)-1]$. For this, agents store an identifier Id_m corresponding to the minimal one he saw. An agent with *id* Id and state $i \in \{0, 1, 2\}$ stores $i \cdot 2^{Id - Id_m}$. When it meets an identifier equals to $Id_m - 1$, it looks for a leader with *id* $Id - 1$ to give it its stored integer (rules would be a bit more precise to handle cases where those leaders have a different Id_m).

Starting from now on, when we say that the population uses its size, we suppose that the counting protocol has been performed and that the protocol uses the chain to access to this information. Once again, the information is globally encoded in the population, and not known by or nor stored in any particular agent.

3.2 Resetting a computation

The computation of the value of n (encoded as above) will be crucial for the following protocols. From now on, we will call *the leader* the (or an) agent with identifier 0 and state $\perp$ even if the previous protocol (computing the size of the population) has not yet finished.

We will now provide a Reset protocol. This protocol has the goal to reach a configuration where the previous protocol is over, where all agents but the leader are in state R, and where the leader knows when this configuration is reached. This protocol will then permit to launch the computation of some other protocols using the chain created and the size of the population computed (i.e. encoded globally in the population) (more explicit proofs can be found in appendix).

▶ Proposition 2 (Reset Protocol). There exists a Reset protocol containing the states F and R such that, once the counting protocol is finished, only one agent will reach state F at some point. As soon as this agent is in state F, all the other agents are in state R.

Proof. The idea of this protocol is to reset each time the leader sees that the counting protocol has not finished yet. The leader tries to turn to state S each other agents, and then turns them to R and count them. When the leader manages to turn m agents (where m is the computed size), it knows that if the counting protocol has finished and the reset protocol is over. ◀

3.3 Counting agents in a given state

We use the previous constructions to create a protocol that can write in its chain[3], with the request of an input symbol $s \in \Sigma$ and an *identifier* Id, the number of agents that started with this *identifier* and this input symbol (The complete proof is in the appendix).

▶ Proposition 3. When we have $f(n) = \Omega(\log n)$ identifiers, if the reset protocol has finished, for all $s \in \Sigma$ and for all $Id \leq f(n)$, there exists a protocol that encodes the number of agents started in the state s_{Id}.

Proof. We cannot use directly the counting protocol here: we cannot store forever this value if the request is done for each s_{Id}, as agents have a finite memory. Because of that, the protocol will need to be sure that the computation is over to move forward and clean the computation. We will use here the fact that the total number of agents is known, by counting the number of agents in the initial state s_k and, at the same time, counting again the whole population. Once we have reached the right total for the population, we know that we have counted all the agents in the requested initial state. ◀

▶ Remark. In other words, if at some moment, the population needs to know the number of agents which started in the state s_{Id}, this is possible.

3.4 Simulating the reading tape

With all these ingredients we will now simulate a tape of a Turing machine. First, we need to define which kind of Turing machines we consider. Basically, we are only stating that from the definitions of the models, only symmetric predicates or data can be processed or computed. Our definitions are an adaptation of the usual models to fit to our inputs.

▶ **Definition 4** (f-Symmetry)**.** A Language $L \in (\Sigma \cup \#)$ is f-symmetric if and only if:

- $\# \notin \Sigma$;
- Words of L are all of the form $w = x_1\#x_2\#\ldots\#x_{f(n)}$, with $|x_1| + |x_2| + \ldots + |x_{f(n)}| = n$ and $\forall i$, $x_i \in \Sigma^+$;
- If, $\forall i$, x'_i is a permutation of x_i, and if $x_1\#x_2\#\ldots\#x_{f(n)} \in L$, then $x'_1\#x'_2\#\ldots\#x'_{f(n)} \in L$;

Each x_i is a non-empty multiset over alphabet Σ.

▶ **Definition 5** ($MNSPACE(S(n))$)**.** Let S be a function $\mathbb{N} \to \mathbb{N}$.
We write $MNSPACE(S(n), f(n))$, or $MNSPACE(S(n))$ when f is unambiguous, for the set of $f-$symmetric languages recognized by Non Deterministic Turing Machines on $O(S(n))$ space.

[3] Recall the concept of the chain defined in page 6. Of course, the idea is that values are encoded in the chain as before.

▶ **Definition 6** (*$SMNSPACE(S(n))$*)**.** We write $SMNSPACE(S(n), f(n))$, or $SMNSPACE(S(n))$ when f is unambiguous, for the set of $f-$symmetric languages recognized by Non Deterministic Turing Machines on $O(S(n))$ space, where the languages are also stable under the permutation of the multisets (i.e. for any permutation σ, $x_1\#x_2\#\ldots\#x_{g(n)} \in L \Leftrightarrow x_{\sigma(1)}\#x_{\sigma(2)}\#\ldots\#x_{\sigma(g(n))} \in L$).

▶ Remark. We have $NSPACE(S(n)) = MNSPACE(S(n), n)$ and $SNSPACE(S(n)) = MNSPACE(S(n), 1)$.

Here is a weaker bound than the one we will obtain. The idea of this proof helps to understand the stronger result.

▶ Proposition 4. Any language in $MNSPACE(\log n, \log n)$ can be recognized by an homonym population protocol with $\log n$ identifiers.

Proof. The main idea of this proof is to use the chain as a tape for a Turing Machine. To simulate the tape of the Turing machine, we store the position where the head of the Turing machine is by memorizing on which multiset the head is (by the corresponding *identifier*) and its relative position inside this multiset: the previous protocol will be used to find out the number of agents with some input symbol in the current multiset, in order to update all these information and simulate the evolution of the Turing machine step by step.

More precisely, let $M \in MNSPACE(\log n, \log n)$. There exists some $k \in \mathbb{N}$ such that M uses at most $k \log n$ bits for each input of size n. To an input $x_1\#x_2\#\ldots\#x_{f(n)}$ we associate the input configuration with, for each $s \in \Sigma$ and for each $i \leq f(n)$, $|x_i|_s$ agents in state k with the identifier $(i-1)$.

The idea is to use the chain as the tape of the Turing Machine. We give k bits to each agent, so that the protocol has a tape of the good length (the chain is of size $\log n$). We just need to simulate the reading of the tape. The protocol starts by counting the population and resetting agents after that.

We assume that symbols on Σ are ordered. Since the language recognized by M is $\log n$-symmetric, we can reorganize the input by permuting the x_is such that the input symbols are ordered.

Here is a way to perform the simulation of reading the tape:

0. The chain contains two counters. The leader also stores an identifier Id and a state s. The first counter stores the total of s_{Id}. The second counter c_2 is the position the reading head: The simulated head is on the c_2th s of x_{Id+1}.

1. At the beginning of the protocol, the population counts the number of S_0, where S_0 is the minimal element of Σ. c_2 is initialized to 1.

2. When the machines needs to go right on the reading tape, c_2 is incremented. If c_2 equals c_1, the protocol looks for the next state s' in the order of Σ, and count the number of s'_{Id}. If this value is 0, then it takes the next one. If s was the last one, then the reading tape will consider to be on a $\#$.
If the reading head was on a $\#$, then it looks for the successor identifier of Id, and counts the number of S_0. If Id was maximal, the machine knows it has reached the end of the input tape.

3. The left movement processus is similar to this one.

This protocol can simulate the writing on a tape and the reading of the input. To simulate the non deterministic part, each time the leader needs to make a choice, it looks for an agent. If its Id is equal to 1, then it does the first choice, otherwise it choses the second possibility.

This protocol simulates M. ◀

▶ **Corollary 7.** *Let f such that $f = \Omega(\log n)$. Any language in $MNSPACE(f(n), f(n))$ can be recognized by an homonym population protocol with $f(n)$ identifiers.*

Proof. We use the same protocol (which is possible as the size of the population can be computed). Since the chain of *identifier*s has a length of $f(n)$, we have access to a tape of size $f(n)$. ◀

3.5 Recognizing polylogarithmic space

▶ Proposition 5. Let f such that $f = \Omega(\log n)$. Let k be a positive integer.

Any language in $MNSPACE\left(\log^k n, f(n)\right)$ can be recognized by a protocol with $f(n)$ identifiers.

Proof. The idea here is that, by combining several *identifier*s together, we get much more identifiers available, increasing the chain and space of computation: if we combine m *identifier*s together, we get $f(n)^m$ possible *identifier*s.

First the population performs the computation of the size of the population. From this, it gets a chain of all the *identifier*s. The leader then creates a counter of m identifiers, initialized at 0^m (seen as the number $0 \dots 0$ written in base $f(n)$). It looks for a $\perp$ agent and gives him its stored m-tuple, then increases its m-tuple. As soon as it has finished (by giving $f(n)^m$ *identifier*s of n *identifier*s, depending on what happens first), the protocol can work on a tape of space $f(n)^m$.

Since $f(n) \geq \epsilon \frac{\log n}{\log \log n}$, there exists some m such that $f(n)^m \geq \log^k n$.

◀

▶ **Theorem 8.** *Let f such that there exists some real $r > 0$ such that we have $f(n) = \Omega(\log^r n)$.*

Any language in $\bigcup_{k \geq 1} MNSPACE(\log^k n, f(n))$ can be recognized by an homonym population protocol with $f(n)$ identifiers.

Proof. We only need to treat the counting protocol when $r < 1$ (the case $r = 1$ is treated in Proposition 5, the case $r > 1$ is a direct corollary of this proposition). Let $l = \lceil \frac{1}{r} \rceil$. We will use a l-tuple for each agent. When agents realize that $f(n)$ might be reached and they need more space, they use the tuple, storing the maximal *identifier* Id_1. If at some point, they realize that a bigger *identifier* Id_2 exists, they just do a translation of the numbers stored in the chain.

With $f(n)^l = \Omega(\log n)$ identifiers and the right basis to write the size, we can be sure to have enough space to compute the size of the population. We can then use previous protocols using $(k \cdot l)$-uples to use the required space.

◀

3.6 Only polylogarithmic space

▶ **Theorem 9.** *Consider a predicate computed by a protocol with $f(n)$ identifiers. Assume that $f(n) = O(\log^l n)$ for some $l \geq 1$. The predicate is in $MNSPACE(\log^k n, f(n))$ for some positive integer k.*

Proof. We need to prove that there exists a Turing Machine that can compute, for any given input x, the output of the protocol P.

From definitions, given some input x, P outputs the output y on input x if and only if there exists a finite sequence $(C_i)_{i \in \mathbb{N}}$, starting from an initial condition C_0 representing x,

that reaches at some finite time j some configuration C_j with interpretation y, and so that any configuration reachable from C_j has also interpretation y.

This latter property can be expressed as a property on the graph of configurations of the protocol, i.e. on the graph whose nodes are configurations of size n, and whose edges corresponds to unique step reachability: one must check the existence of a path from C_0 to some C_j with interpretation y so that there is no path from C_j to some other C' with interpretation different from y.

Such a problem can be solved in $NSPACE(\log N)$ where N denotes the number of nodes of this graph. Indeed, guessing a path from C_0 to some C_j can easily be done in $NSPACE(\log N)$ by guessing intermediate nodes (configurations) between C_0 and C_j. There remain to see that testing if there is no path from C_j to some other C' with interpretation different from y can also be done in $NSPACE(\log N)$ to conclude.

But observe that testing if there is a path from C_j to some other C' with interpretation different from y is clearly in $NSPACE(\log N)$ by guessing C'. From Immerman-Szelepcsnyi's Theorem [**?**, **?**] we know that $NSPACE(\log N) = coNSPACE(\log N)$. Hence, testing if there is no path from C_j to some other C' with interpretation different from y is indeed also in $NSPACE(\log N)$.

It remains now to evaluate N: For a given identifier i, an agent encodes basically some basic state $b \in B$, and d identifiers $u_1, u_2, \ldots, u_d$. There are at most n agents in a given state $(i, b, u_1, u_2, \ldots, u_d)$. Hence $N = O(n^{|B| \cdot f(n)^{d+1}})$. In other words, the algorithm above in $NSPACE(\log N)$ is hence basically in $MNSPACE((|B| \cdot f(n)^{d+1}) \log n, f(n)) \subset MNSPACE(\log^k n, f(n))$ for some k. ◀

▶ **Theorem 10.** *Let f such that for some r, we have $f(n) = \Omega(\log^r n)$. The set of functions computable by homonym population protocols with $f(n)$ identifiers corresponds exactly to $\bigcup_{k\geq 1} MNSPACE(\log^k n, f(n))$.*

3.7 When we have $\sqrt[k]{n}$ identifiers

One can go from $\sqrt[k]{n}$ to a space of computation equivalent to the case where $f(n) = n$: We just need to use a k-tuple of identifiers.

▶ **Theorem 11** ($n^{1/k}$ Identifiers)**.** *Let f such that there exists some $k \in \mathbb{N}$ such that $f(n) \geq n^{1/k}$. The set of functions computable by homonym population protocols with $f(n)$ identifiers corresponds exactly to $MNSPACE(n \log n, f(n))$.*

▶ Remark. This result does not need the two restrictions of knowing if an *id* is equal to 0 or if two *id*s are consecutive. The result holds when U is chosen arbitrarily and when the restrictions over the rules are those in [**?**].

4 Passively Mobile Machines

We now show how previous constructions improve the results about the model from [**?**]:

▶ **Definition 12** ($PMSPACE(S(n))$)**.** [**?**] Let S be a function. We write $PMSPACE(S(n))$ for the set of languages recognized by population protocols where each agent has a Turing Machine with a tape of size a least $S(n)$.

▶ **Theorem 13.** $PMSPACE(\log \log n) = \bigcup_{k\geq 1} SNSPACE(log^k n)$.

Proof. **1.** $\bigcup_{k\geq 1} SNSPACE(log^k n) \subset PMSPACE(\log\log n)$.

The idea of this proof is quite simple: Let $M \in SNSPACE(log^k n)$. We can notice that $SNSPACE(log^k n) \subset MNSPACE(log^k n, \log n)$. From Theorem 10, there is a population protocol computing M. We will simulate it. With space $O(\log\log n)$, we can simulate a population protocol with $O(2^{\log\log n}) = O(\log n)$ identifiers.

To create $\log n$ identifiers, we adapt a bit the counting protocol. At the beginning, each agent has the identifier 0. When two agents with the same *identifier* meet, if each one contains the integer 1, the first switch its integer to 0, the other increases its own *identifier*.

We then just need to simulate the behavior of each agent as if they have started with their created *identifier*. It requires a space of size $|B| + (d+1)\log\log n$ plus some constant, which is enough.

2. $PMSPACE(\log\log n) \subset \bigcup_{k\geq 1} SNSPACE(log^k n)$: The proof is similar to the one of Theorem 9.

◀

With a similar proof, we can get the following result that gives a good clue for the gap between $\log\log n$ and $\log n$:

▶ **Corollary 14.** *Let f such that $f(n) = \Omega(\log\log n)$ and $f(n) = o(\log n)$.*
$SNSPACE(2^{f(n)} f(n)) \subset PMSPACE(f(n)) \subset SNSPACE(2^{f(n)} \log n)$.

5 Summary

From the model given by Guerraoui and Ruppert [**?**], we introduced a hierarchy according to the number of accessible identifiers in the population:

- With a constant number of the identifiers, the existence of identifiers is useless.
- With $\Theta(\log^r)$ identifiers, homonym population protocols can exactly recognize any language in $\bigcup_{k\in\mathbb{N}} MNSPACE\left(\log^k n\right)$.
- Homonym population protocols with $\Theta(\sqrt[r]{n})$ identifiers have same power that homonym population protocols with n identifiers.

It remains an open question: is the knowledge of consecutive values of two identifiers crucial or not? Our guess is that it is essential.

Chatzigiannakis *et al.* [**?**] started a hierarchy over protocols depending on how much space of computation we provide to each agent. The paper left an open question on the gap between $o(\log\log n)$ and $O(\log n)$. We provided an answer, stating that with $\Theta(\log\log n)$ space, we compute exactly $\bigcup_{k\in\mathbb{N}} SNSPACE\left(\log^k n\right)$.

It remains the gap between $O(\log\log n)$ and $O(\log n)$, where we currently just have the following bounds: $SNSPACE(2^{f(n)} f(n)) \subset PMSPACE(f(n)) \subset SNSPACE(2^{f(n)} \log n)$.

A Proof of Proposition 2

Each time one of the interaction in the counting protocol happens, one of the two interacting agents looks for a leader (their might be several ones) and tells it to start back the reset protocol. The leader at this moment goes to state G.

Each agent has a triplet of states. The first one corresponds to the counting protocol. The second one corresponds to the reset part. The third one is here to perform another counting protocol. The leader's goal is to turn all the agents in state R, counting these agents at the same time with the third state. When it manages to turn the same number of agents that the number counted by the protocol of Proposition 1, it knows that its work is over.

The protocol works in 6 steps. Each step can be deduced from the current state of the leader, being G for step 1, H for step 2...:

1. The leader turns each agent into state S with the rule $G_0 \quad q_k \to G_0 \quad S_k$. We add the rule $q_k \quad G_0 \to S_k \quad H_0$, which permits to go to the second step of the protocol.
2. The leader looks for an agent. If it meets one in state S, it turns the agent in state R, then the protocol goes to step 3. If it meets an agent in state $q \neq S$, it goes to step 4.
3. The leader increments the third counting protocol. During the incrementation (the leader goes throw the chain to perform it), it also checks bit by bit if the first and second counted numbers are the same. If yes, the protocol goes to step 5, else it goes back to step 2.
4. The leader restarts its count, putting back each bit to 0 in the 3rd state of the agents of the chain. It also resets its bit to 1. Then, the protocol goes back to step 1.
5. The leader is sure that all the other agents in the population have been turned from S to R. It changes its own state to F and stops the Reset protocol.

By fairness, when the counting protocol is finished, we can be sure that the protocol will be able to reach the 5th step. When the 5th step is reached, we are sure that at least $(n-1)$ agents are in state R, as the leader counted the agents it turned to R. We can conclude that the protocol did what was expected.

B Proof of Proposition 3

We suppose that the population is already reseted to the state N.

We do not give here the exact protocol, only how it works:

0. The counting chain will use 4 bits: one for the counting protocol (that is supposed to be already finished), one equals to one if and only if the leader has already counted it during this protocol, one for the counting of agents in the input s_{Id}, and one for computing again the total. The last element is here to know when the leader has met every agent to check if they have started in state s_{Id} or not.

   When we say the second counter, it is the one using the 3rd bit, the third counter using the 4th bit.
1. The leader looks for an agent it has not recounted again (i.e. with its second bit equals to 0). When it meets one, its switches its second bit, it looks if its input was s_{Id} or not. If it is, it increments the second and the third counter, otherwise it increments only the third.
2. The leader then looks if the first and the third counter are equal. If not, it goes back to step 1, if yes the computation is over.

Since the counting protocol is over (if not, the population will be reseted again and again until the counting is over), the size is known. With that, we are sure to have counted each agent started in state s_{Id}, as the leader must have seen each agent in this protocol before finishing it.

The assumption that each agent started in state N is not a big deal, as the leader may look for each agent, turn them into state N and count in parallel how many agents it has turned into state N.